

Stationary Mixing Bandits

¹, Julien Audiffren¹, and Liva Ralaivola²

¹CMLA, ENS Cachan, Cachan, France

²QARMA, Aix-Marseille Universit, CNRS, LIF, Marseille, France

Abstract

We study the bandit problem where arms are associated with stationary φ -mixing processes and where rewards are therefore dependent: the question that arises from this setting is that of recovering some independence by ignoring the value of some rewards. As we shall see, the bandit problem we tackle requires us to address the exploration/exploitation/independence trade-off. To do so, we provide a UCB strategy together with a general regret analysis for the case where the size of the independence blocks (the ignored rewards) is fixed and we go a step beyond by providing an algorithm that is able to compute the size of the independence blocks from the data. Finally, we give an analysis of our bandit problem in the restless case, i.e., in the situation where the time counters for all mixing processes simultaneously evolve.

1 Introduction

Bandit with mixing arms. The bandit problem consists in an agent who has to choose at each step between K arms. A stochastic process is associated to each arm, and pulling an arm produces a reward which is the realization of the corresponding stochastic process. The objective of the agent is to maximize its long term reward. It is classically assumed that the stochastic process associated to each arm is a sequence of independently and identically distributed (i.i.d) random variables (see, e.g. [1]). In that case, the challenge the agent has to face is the well-known exploration/exploitation problem: she has to simultaneously make sure that she collects information from all arms to try to identify the most rewarding ones —this is *exploration*— and to maximize the rewards along the sequence of pulls she performs —this is *exploitation*. Many algorithms have been proposed to solve this trade-off between exploration and exploitation [2, 3, 4, 1]. We propose to go a step further than the i.i.d setting and to work in the situation where the process associated with each arm is a stationary φ -mixing process and the rewards are thus dependent from one another, but with a strength of dependence that weakens over time. From an application point of view, this is a reasonable dependence structure: if a user clicks on some ad (a typical use of bandit algorithms) at some point in time, it is very unlikely that she will click again on this ad in the near future. As it shall appear in the sequel, working with such dependent observations poses the question of how informative are some of the rewards with respect to the value of an arm since, because of the dependencies and the high correlation between close-by (in time) rewards, they might not reflect the true 'value' of the arms. However, as the dependencies weaken over time, some kind of independence might be recovered if rewards are ignored. This actually requires to deal with a new trade-off exploration/exploitation/independence that need be precisely handled.

Rested and Restless case. A closely related setup that addresses the bandit problem with dependent rewards is when they are distributed according to Markov processes, such as Markov chains and Markov decision process (MDP) [5, 6], where the dependencies between rewards are of bounded range, which is what distinguishes those works with ours. Contributions in this area study two settings, that we will analyze as well: the *rested* case, where the process attached to an arm evolves only when the arm is pulled, and the *restless* case, where all processes simultaneously evolve at each time step.

Contributions and structure of the paper. We define the notion of a φ -mixing bandit and its regret (Section 2), we provide a general analysis and an algorithm to solve the rested case where the size of independence blocks is fixed (Section 3). We provide another approach where these sizes are computed from the data by introducing another algorithm (Section 4). Finally, in Section 5, we provide an algorithm and a regret analysis to deal with the restless case.

2 Overview of the problem

Let $(\omega, \mathcal{F}, \mathbb{P})$ be a probability space. We recall the definitions of stationary and of φ -mixing processes:

Definition 1 (Stationarity). A sequence of random variables $\underline{X} = \{X_t\}_{t=-\infty}^{+\infty}$ is *stationary* if, for any t and nonnegative integer m and s , the random subsequences (X_t, \dots, X_{t+m}) and $(X_{t+s}, \dots, X_{t+m+s})$ are identically distributed.

Definition 2 (φ -mixing process). Let $\underline{X} = \{X_t\}_{t=-\infty}^{+\infty}$ be a stationary sequence of random variables. For any $i, j \in \mathbb{Z} \cup \{-\infty, +\infty\}$, let σ_i^j denote the σ -algebra generated by the random variables X_t , $i \leq t \leq j$. Then, for any positive integer n , the φ -mixing coefficient $\varphi(n)$ of the stochastic process \underline{X} is defined as

$$\varphi(n) = \sup_{t, A \in \sigma_{t-k}^{+\infty}, B \in \sigma_{-t}^t} |\mathbb{P}[A|B] - \mathbb{P}[A]|. \quad (1)$$

\underline{X} is said to be φ -mixing if $\varphi(n) \rightarrow 0$ as $n \rightarrow \infty$.

We are interested in the problem of sampling from a K -armed φ -mixing bandit. In our setting, pulling arm k at time t provides the agent with a realization of the random variable $X_{\tau_k(t)}^k$, where $\tau_k(t) = t$ in the restless case and $\tau_k(t)$ is the number of times arm k was pulled in the rested case, and where the family $(X_t^k)_{t \geq 1}$ satisfies the following hypotheses :

1. $\forall k, (X_t^k)_{t \in \mathbb{Z}}$ is stationary;
2. the sequences $(X_t^k)_{t \in \mathbb{Z}}$ are φ -mixing;
3. each X_1^k takes values in a discrete finite set.

This setting assumes the possibility of long-term dependencies between the rewards output by the arms. It is important to note that by definition of the φ -mixing processes, the amount of dependence decreases with time. Hence, as evoked earlier, in order to choose which arm to pull, the agent is forced to address the exploration/exploitation/independence trade-off where *independence* may be partially recovered by ignoring some rewards so as to make computations on data that are distant in time, i.e. data that are not too correlated (thanks to the mixing property).

It is critical to note that unlike in the i.i.d. framework, Hoeffding inequality cannot be applied in this setting, thus the widely used upper confidence bound (UCB) algorithms cannot be used here. In the case of stationary φ -mixing distributions, we have the following theorem from [7].

Theorem 1 ([7, 8]). *Let $\psi : \mathcal{U}^m \rightarrow \mathbb{R}$ be a function defined over a countable space \mathcal{U} , and \underline{X} be a stationary φ mixing process. If ψ is l -Lipschitz with respect to the Hamming metric for some $l > 0$, then the following holds for all $t > 0$:*

$$\mathbb{P}_{\underline{X}}[|\psi(\underline{X}) - \mathbb{E}\psi(\underline{X})| > t] \leq 2 \exp \left[-\frac{t^2}{2ml^2 \|\Lambda_m\|_{\infty}^2} \right], \quad (2)$$

where $\|\Lambda_m\|_{\infty} \leq 1 + 2 \sum_{k=1}^m \varphi(k)$.

In the following, we consider a more general framework than the one usually encountered in the bandit literature. Instead of looking at cumulative gains, we look at rewards computed according to Lipschitz functions meeting the requirements of the concentration inequality stated in Theorem 1.

More precisely, we suppose that we have K known families $(\psi_t^k)_{t \geq 1}$ $k = 1, \dots, K$ of functions such that for every non-negative integer m :

1. $\psi_m^k : \mathcal{U}^m \rightarrow \mathbb{R}$ accounts for the reward associated with m consecutive outcomes of arm k ;
2. ψ_m^k is 1-Lipchitz with respect to the Hamming metric.

In the following sections, we use m to identify the reward functions $(\psi_m^k)_{1 \leq k \leq K}$ we want to optimize, b to refer to the size of the independence blocks and we consider blocs of s trials. We study three different scenarios. In the next section, we will present a general algorithm and regret analysis in the rested case where m is fixed and $s \doteq m + b$. In Section 4, with an additional hypothesis on the φ -mixing processes, we will take another approach to the rested case when $m + b \mid s$ (i.e. s is a multiple of $m + b$) by including the notion of dependency into the regret. Finally, in Section 5, we present a general algorithm and regret analysis for the restless case.

3 Mixing Bandits in the Rested Case

Regret analysis. Here, we are going to analyze the situation where m and b are *fixed*. Our goal is to show that a simple algorithm derived from UCB that works by making blocks $s \doteq m + b$ consecutive trials on each arm has low regret, with a notion of regret that we define in the sequel. The running time index is therefore of the form of st , with t the number of times arm selection has been performed, and the sequence accessed to are such as $(X_{st}^k, \dots, X_{(s+1)t-1}^k)$. The reward that is accessed to at time t (with a slight abuse of notation that makes us use t as the time index) when pulling arm k does not make use of the full information provided by this sequence but instead is $\psi_m^k(\underline{X}_{t,s,b}^k)$ where

$$\underline{X}_{t,s,b}^k \doteq (X_{st}^k, \dots, X_{(s+1)t-b-1}^k); \quad (3)$$

this means that only the first m points from $(X_{st}, \dots, X_{(s+1)t-1})$ are taken advantage of.

Given τ , the total number of trials of s -blocks, the regret that we are going to work with is

$$\mathcal{R} \doteq \tau \mu_{\psi,m}^* - \sum_{k=1}^K \mathbb{E} \tau_k(\tau) \mu_{\psi,m}^k \quad (4)$$

where: $\mu_{\psi,m}^* \doteq \max_{k=1,\dots,K} \mu_{\psi,m}^k$, $\mu_{\psi,m}^k \doteq \mathbb{E}_{X_1^k, \dots, X_m^k} \psi_m^k(X_1^k, \dots, X_m^k)$ and $\tau_k(t)$ is the number of times arm k has been chosen given that a total (i.e. over all arms) of t pulls of s -blocks have been performed.

The arm selection procedure of the algorithm that we propose and dub Block-UCB, is depicted in Algorithm 1, where the function Λ_k is defined as

$$\Lambda_k(t) \doteq 1 + 2 \sum_{r=1}^t \varphi^k(rb + (r-1)m). \quad (5)$$

It is possible to show that Block-UCB has the following regret.

Theorem 2 (Regret of Block-UCB). *The regret of Block-UCB is bounded by*

$$\sum_{k: \mu_{\psi,m}^k < \mu_{\psi,m}^*} \left(u_k \Delta_k + \frac{1}{\alpha - 2} \right),$$

where the u_k 's are the solutions of the problems

$$u_k \Delta_k^2 - 8\alpha \Lambda_k^2(u_k) \log \tau = 0, \quad k = 1, \dots, K. \quad (6)$$

The result of Theorem 2 hinges on the derivation of a concentration inequality for each arm k that relates the random variable $\frac{1}{\tau} \sum_{r=0}^{\tau-1} \psi_m^k(\underline{X}_{r,s,b}^k)$ to $\mu_{\psi,m}^k$. To establish this concentration inequality, we study the random variables $\Gamma_b^k(\underline{X}_{0,s,b}^k, \dots, \underline{X}_{\tau-1,s,b}^k)$, defined for $b \geq 0$ as

$$\Gamma_b^k(\underline{X}_{0,s,b}^k, \dots, \underline{X}_{\tau-1,s,b}^k) \doteq \frac{1}{\tau} \sum_{r=0}^{\tau-1} \psi_m^k(X_{rs}^k, \dots, X_{(r+1)s-b-1}^k) - \mu_{\psi,m}^k. \quad (7)$$

The concentration inequality that we are going to use to prove our regret bound is the following:

Algorithm 1 Main iteration of Block-UCB

Choose arm I_t as

$$I_t \in \operatorname{argmax}_k \frac{1}{t} \sum_{r=0}^{t-1} \psi_m^k(\underline{X}_{r,s,b}^k) + \Lambda_k(\tau_k(t-1)) \sqrt{\frac{2\alpha \log t}{\tau_k(t-1)}}.$$

Theorem 3. For all τ, k, b , and assuming that ψ_m^k takes value in $[0; 1]$:

$$\mathbb{P}\left(\left|\Gamma_b^k(\underline{X}_{0,s,b}^k, \dots, \underline{X}_{\tau-1,s,b}^k)\right| \geq \varepsilon\right) \leq \exp\left(-\frac{\tau \varepsilon^2}{2\Lambda_k^2(\tau)}\right), \quad (8)$$

where Λ_k is defined in (5).

Proof. We make the proof for some arm k . We also assume that s and b are fixed and to lighten the notation, we drop the dependence on these variables when no confusion is possible: we use \underline{X}_r (resp. Γ) for $\underline{X}_{r,s,b}^k$ (resp. Γ_b^k). The proof hinges on the fact that since $(X_t)_{t \geq 0}$ is a stationary mixing sequence with mixing coefficients $(\varphi(t))_{t \geq 1}$, $(\underline{X}_r)_{r \geq 0}$ is a stationary mixing sequence with mixing coefficients $(\varphi(q))_{q \geq 1}$ such that $\varphi(q) \doteq \varphi(qb + (q-1)m)$ (see Proposition 5, Appendix). To obtain the targeted result, we make use of the concentration inequality of Theorem 1 with the function Γ to get (8) and we exploit the observation that $\mathbb{E}\Gamma_b^k = 0$, which results from the stationarity of $(X_t)_{t \geq 0}$.

Let q be an integer in $\{0, \dots, r-1\}$, $\underline{x}_0, \dots, \underline{x}_{r-1}$ and \underline{x}'_q blocks from \mathcal{U}^{s-b} . Then:

$$\left|\Gamma(\underline{x}_0, \dots, \underline{x}_q, \dots, \underline{x}_{r-1}) - \Gamma(\underline{x}_0, \dots, \underline{x}'_q, \dots, \underline{x}_{r-1})\right| = \left|\frac{1}{\tau}(\psi_{s-b}(\underline{x}_q) - \psi_{s-b}(\underline{x}'_q))\right| \leq \frac{1}{\tau},$$

which comes from the range of ψ_{s-b} being $[0; 1]$. Γ is therefore $1/\tau$ -Lipschitz with respect to the Hamming metric, which, combined with $(\underline{X}_r)_{r \geq 0}$ being a φ -mixing sequence, gives (8). \square

The proof of the previous theorem uses the following more general result, that is of independent interest.

Theorem 4 (General Regret). Suppose that the arms we work with are such that

$$\forall k, \mathbb{P}(|\hat{\mu}_\tau^k - \mu^k| \geq \varepsilon) \exp(-\theta_k(\tau)\gamma_k(\varepsilon)), \quad (9)$$

where τ is the number of data the empirical mean $\hat{\mu}_\tau^k$ is computed on, and θ_k and ε_k are increasing functions defined on $(0; +\infty]$.

Consider the regret defined by

$$R \doteq \tau \mu^* - \sum_{k=1}^K \mathbb{E} \tau_k(\tau) \mu_k$$

where $\tau_k(t)$ is the number of times a (suboptimal) arm k has been chosen up to time t .

The (α, θ, γ) -UCB that chooses at iteration t an arm I_t according to

$$I_t \in \operatorname{argmax}_k \hat{\mu}_{\tau_k(t-1)}^k + \gamma_k^{-1}\left(\frac{\alpha}{\theta_k(\tau_k(t-1))} \log t\right)$$

has regret bounded by:

$$\sum_{k: \mu_i < \mu^*} \left(\left\lceil \theta_k^{-1}\left(\frac{\alpha \log \tau}{\gamma_k(\Delta_k/2)}\right) \right\rceil \Delta_k + \frac{1}{\alpha - 2} \right). \quad (10)$$

Proof. Note that Theorem 2 is a consequence of this theorem with $\theta_k(s) = s/\Lambda_k^2(s)$ and $\gamma_k(\varepsilon) = \varepsilon^2/2$.

The proof use the standard technique to prove the regret of UCB-like algorithms. Namely, at iteration t , if $I_t = i$ for i not optimal, then one of the following events $\mathcal{E}_1^*(t), \mathcal{E}_2(i, t), \mathcal{E}_3(i, t)$ must occur

$$\mathcal{E}_1^*(t) \doteq \left\{ \hat{\mu}^* < \mu^* - \gamma_*^{-1} \left(\frac{\alpha}{\theta_*(\tau^*(t-1))} \log t \right) \right\}, \quad (11)$$

$$\mathcal{E}_2(i, t) \doteq \left\{ \mu^i \leq \hat{\mu}^i - \gamma_i^{-1} \left(\frac{\alpha}{\theta_i(\tau_i(t-1))} \log t \right) \right\}, \quad (12)$$

$$\mathcal{E}_3(i, t) \doteq \left\{ \tau_i(t-1) \leq \theta_i^{-1} \left(\frac{\alpha \log \tau}{\gamma_i(\Delta_i/2)} \right) \right\}. \quad (13)$$

Indeed, if none of the events occurs then (using $\Delta_i = \mu^* - \mu_i$)

$$\hat{\mu}^* + \gamma_*^{-1} \left(\frac{\alpha}{\theta_*(\tau^*(t-1))} \log t \right) \stackrel{(11)}{\geq} \mu_i + \Delta_i \stackrel{(12)}{>} \mu_i + 2\gamma_i^{-1} \left(\frac{\alpha}{\theta_i(\tau_i(t-1))} \log \tau \right) \stackrel{(13)}{>} \hat{\mu}_i + \gamma_i^{-1} \left(\frac{\alpha}{\theta_i(\tau_i(t-1))} \log t \right)$$

where we have used that $t \mapsto \gamma_i^{-1}(a \log t)$ is an increasing function of t on $[1; \infty)$ whenever $a > 0$. This implies that $I_t \neq i$, which contradicts our working hypothesis.

If we let u be defined as:

$$u \doteq \left\lceil \theta_i^{-1} \left(\frac{\alpha \log \tau}{\gamma_i(\Delta_i/2)} \right) \right\rceil,$$

then, for i suboptimal, we have the following

$$\begin{aligned} \mathbb{E} \tau_i(\tau) &= \sum_{t=1}^{\tau} \mathbb{E} \mathbb{1}_{[I_t=i]} \leq u + \sum_{t=u+1}^{\tau} \mathbb{E} \mathbb{1}_{[I_t=i \wedge \neg \mathcal{E}_3(i, t)]} \leq u + \sum_{t=u+1}^{\tau} \mathbb{E} \mathbb{1}_{[\mathcal{E}_1^*(t) \vee \mathcal{E}_2(i, t)]} \\ &\leq u + \sum_{t=u+1}^{\tau} [\mathbb{P}(\exists t : \mathcal{E}_1^*(t)) + \mathbb{P}(\exists t : \mathcal{E}_2(i, t))] \end{aligned}$$

Using the union bound and Equation (9), both probabilities $\mathbb{P}(\exists t : \mathcal{E}_1^*(t))$, $\mathbb{P}(\exists t : \mathcal{E}_2(i, t))$ can be bounded from above by $1/t^\alpha$. Standard calculations allow us to get desired result (10). \square

Discussion. Some observations must be made regarding the result of Theorem 2. First, as we used Theorem 4 to prove our regret bound, it is necessary for the result to hold for the functions $\theta_k \doteq \tau / \Lambda_k^2(\tau)$ to be increasing. In addition, the regret only makes sense if it is bounded, i.e. if the $u_k \Delta_k$ are bounded. Finally, if these conditions hold, it might be interesting to find, for some fixed horizon τ , the value of b that minimizes the regret. We now depicts common settings for the φ^k that yield instructive results and that build upon the previous remarks. For the sake of conciseness, we will assume that all $\varphi^1 = \dots = \varphi^K$ and we use φ to refer to the mixing coefficients.

Independent case. if $\varphi = 0$, i.e. we are in the independent case, and the θ_k 's are naturally increasing. In addition, it is straightforward to observe that the best use of the data is achieved for $b = 0$, i.e. each and every reward is used to estimate the quality of an arm.

Case $\Lambda < +\infty$. In that case, we are again back to a situation almost similar to the usual independent case. The θ_k 's are increasing, the u_k are well-defined and the regret as the usual $O(\sum_k \log \tau / \Delta_k)$ form.

Algebraically mixing case. Here, $\varphi(t) = \varphi_0 t^{-p}$ for $p > 1$, and a few calculations give

$$\begin{aligned} \Lambda(\tau) &= 1 + 2\varphi_0 \sum_{r=1}^{\tau} \frac{1}{(rs - m)^p} \leq 1 + 2\varphi_0 \left(1 + \int_1^{\tau} \frac{1}{(rs - m)^p} dr \right) \\ &= 1 + 2\varphi_0 + \frac{2\varphi_0}{s(p-1)} \left(\frac{1}{b^{p-1}} - \frac{1}{(\tau s - m)^{p-1}} \right) \end{aligned}$$

Using this upper bound to find the u_k s as in (6) and to solve for b so that this bound is minimized provides a way to find a data-dependent b . Another (coarser) way to look at the algebraically mixing situation is not to optimize for b and to consider that it is a particular case of the previous case, since $\sum_t \varphi(t) < \infty$, i.e. $\Lambda < +\infty$. This assumption is made in the rest of the paper.

4 $m + b \mid s$: Expressing the Independence Trade-Off in the Regret

This section introduces another way of encoding the trade-off between exploration, exploitation and independence. As before, we consider sequences of s trials, but among the s results, we seek to optimize the number m of results we use to update the empirical value of the arms and the number b of results we ignore in order to improve the independence between the considered realization of our random variables.

In addition to the case $s = m + b$, we also consider the situation where $\beta(m + b) = s$ with $\beta \in \mathbb{N}$ and $\beta > 1$. The sequence of s trials can then be interpreted as β successive sequences of $m + b$ trials, and the value of this particular (m, b) distribution is thus multiplied by β .

4.1 Hypotheses and Regrets

In this section we make the following additional assumption on the φ -mixing processes $(X_t^k)_{t \geq 0}$:

$$\forall 1 \leq k \leq K, \quad \forall b \in \mathbb{N}, \quad M_k(b) \doteq 1 + \sum_{i \geq 1} \varphi^k(b(i+1)) < +\infty.$$

Note that this is equivalent to the widely used assumption that the $\varphi^k(i)$ are summable over i (see for instance the case of algebraically mixing sequence mentioned before). Also, note that $M_k(1)$ is an upper bound of $(\Lambda_m^k)_m$ which appears in Theorem 1, and $M_k(\cdot)$ is a decreasing function such that $M_k(b) \geq 1$.

The setting is the following: at each step, the agent pulls an arm s times, and has to choose how to split those s elements between a meaningful part of m elements, used to update empirical values of the arm, and the non-significant part of b element, used to strengthen the independence between the variables. For each such combination $m + b \mid s$ and for each arm k , we define the value of the combination (m, b, k) as:

$$\nu_{m,b}^k \doteq \frac{\beta_{m,b}}{M_k(b)} \mu_m^k, \quad (14)$$

where $\mu_m^k \doteq \mathbb{E}_{X_1^k, \dots, X_m^k} \psi_m^k(X_1^k, \dots, X_m^k)$ and $\beta_{m,b} = s/(m + b)$. This value explicitly shows the trade-off between independence (through $M_k(b)$) and exploitation (through the $\mu_{m,b}^k$). The value of an arm k is then defined as the maximum value of the possible combination (m, b, k) , for $m + b \mid s$.

With this in mind, we define the regret \mathcal{R} at time t as :

$$\mathcal{R} \doteq \sum_{k=1}^K \mathbb{E}(T(k)) \left(\nu_{m^*, b^*}^{k^*} - \nu_{m_k^*, b_k^*}^k \right) \quad (15)$$

where $(k^*, m^*, b^*) = \arg \max_{(k, m, b)} \nu_{m, b}^k$, $(m_k^*, b_k^*) = \arg \max_{(m, b)} \nu_{m, b}^k$, and $T(k)$ denotes the number of times arm k was pulled.

It is important to note that one of the main difference between (15) and the classical formulation of regret from a multi arm bandit in the i.i.d. case is that in our setting, we are comparing the value of the best combination of the best arm with the value of the best combination of the pulled arm.

4.2 Concentration inequality and algorithm

We now introduce a concentration inequality tailored for the $\nu_{m,b}^k$ of (14).

Proposition 1. *Let $1 \leq k \leq K$, $1 \leq m \leq s$, $b = m - s$, $n \in \mathbb{N}^*$, $\psi_m^k : \mathcal{U}^m \rightarrow \mathbb{R}$ be a 1-Lipschitz with respect to the Hamming metric function defined over a countable space \mathcal{U} . Suppose that the $\varphi^k(i)$ are summable over i , and let us define $M_k(b) = 1 + \sum_{i \in \mathbb{N}^*} \varphi^k(i(b+1))$, and $\zeta_m^s : \mathbb{N} \mapsto \mathbb{N}$, $\zeta_m^s(t) = t + b \lfloor (t-1)/m \rfloor$. Then the following holds for all $t > 0$:*

$$\mathbb{P} \left[\frac{1}{M^k(b)} \left| \frac{1}{n} \sum_{i=0}^{n-1} \psi_m^k \left(X_{\zeta_m^s(im+1)}^k, \dots, X_{\zeta_m^s((i+1)m)}^k \right) - \mu_m^k \right| > t \right] \leq 2 \exp \left[-\frac{nt^2}{2} \right]. \quad (16)$$

Algorithm 2 Block-UCB with parameters s, α fixed

$t \leftarrow 0, \hat{\psi}_{m,b,0}^k \leftarrow 0, k = 1, \dots, K, m+b \mid s, \tau_{k,0} \leftarrow 0, k = 1, \dots, K$
for $t = 1, \dots, \tau$ **do**

Select arm \hat{k} with $\hat{k} \in \operatorname{argmax}_k \max_{m+b \mid s} \frac{s}{m+b} \left(\hat{\psi}_{m,b,t-1}^k + \sqrt{\frac{2\alpha(m+b)\log(t)}{s\tau_{k,t-1}}} \right)$

Update the block counts $\tau_{k,t} \leftarrow \tau_{k,t-1} + \delta_{k\hat{k}}, \forall k$

Compute the values of the $\hat{\psi}_{m,b,t}^k, \forall m+b \mid s$

$$\begin{aligned} \hat{\psi}_{m,b,t}^{\hat{k}} &\leftarrow \frac{1}{M_k(b)\beta\tau_{\hat{k},t}} \sum_{r=0}^{\beta(\tau_{\hat{k},t}-1)} \psi_m^{\hat{k}} \left(X_{r(m+b)}^{\hat{k}}, \dots, X_{(m+b)r+m-1}^{\hat{k}} \right) \\ \hat{\psi}_{m,t}^k &\leftarrow \hat{\psi}_{m,t-1}^k, \text{ for } k \neq \hat{k} \end{aligned}$$

end for

Proof. This proposition naturally follows from the proof of Theorem 1 with the function $\phi_n^k = \frac{1}{n} \sum_{i=0}^{n-1} \psi_m^k(X_{\zeta_m^s(im+1)}, \dots, X_{\zeta_m^s((i+1)m)})$ which is $1/n$ -Lipschitz with respect to the Hamming metric, and Proposition 5 (see Appendix). \square

Note that the upper bound on the probability that appears in (16) is uniform over (k, m, b) . This is crucial to define our algorithm and to analyze its regret (more details in the next subsection). We now introduce Algorithm 2 that is designed for the particular setting of φ -mixing bandit problem just described. First, note that since the pair (m, b) (with $m+b \mid s$) which gives the best result for each arm is unknown, the algorithm needs to compute an empirical estimator for each of these combinations for each arm. In other words, the algorithm needs to efficiently and simultaneously learn both the best combination and the best arm.

4.3 Regret Analysis

In this subsection we provide an upper bound for the regret of Algorithm 2.

Proposition 2. Let \mathcal{R} be the regret as defined in (15) and let $\eta : \mathbb{N} \mapsto \mathbb{N}, \eta(s) \doteq \sum_{i=1}^s i \mathbb{1}_{i \mid s}$. Then, with $\Delta_i = \nu_{m^*, b^*}^{k^*} - \nu_{m_i^*, b_i^*}^i$

$$\mathcal{R} \leq \sum_{1 \leq i \leq K} \left((1 + \eta(s)) \Delta_i + \frac{8\alpha s \log(t)}{\Delta_i} \right).$$

Proof. The main difference with the standard technique to proving regret bounds comes from the fact that the value of each arm is the maximum of its coordinate: as suchs we have to consider the following event

$$\begin{aligned} \psi_t^{k^*} &\leq \nu_{m^*, b^*}^{k^*} \\ \exists m+b \mid s, \quad \nu_{m_k^*, b_k^*}^k &\leq \frac{s}{m+b} \left(\hat{\psi}_{m,b,t-1}^k - \sqrt{\frac{2\alpha(m+b)\log(t)}{s\tau_{k,t-1}}} \right) \\ \exists m+b \mid s, \quad 2\sqrt{\frac{2\alpha(m+b)\log(t)}{s\tau_{k,t-1}}} &\geq \nu_{m^*, b^*}^{k^*} - \nu_{m_k^*, b_k^*}^k \end{aligned}$$

By carefully using the property of the maximum, the result can be recovered. All the details of the proof can be found in the supplementary material. \square

We have just seen another approach to the rested mixing bandit. By assuming the summability of the φ_k , and by properly defining the value of an arm, we were able using Algorithm 2 to address the case where m and b are no longer fixed, but must be computed from the data by the agent.

It is interesting to note that the upper bound in Proposition 2 differs from the usual bound in the classical dependence-free setting by two multiplicative constants: $\eta(s)$, which encodes the total number of combination of pair (m, b) such that $m + b \mid s$, and s , which is in fact used as an upper bound for $s/(m + b)$.

5 Restless φ mixing bandits

In this section, we provide an analysis for the restless φ -mixing bandit. Recall that contrarily to the rested case studied previously, the stochastic processes associated to each arm evolves regardless of the actions of the agent. This difference is of paramount importance in the φ -mixing setting. Indeed, in the rested case, the agent was bound to ignore some realizations obtained from an arm to enforce the independence and therefore the accuracy of its predictor. In the restless case, instead of pulling an arm to no avail, an agent willing to increase the independence of the realization of an arm k can pull another arm $k' \neq k$ to gather information about k' while enforcing the independence of k . This idea is central to this section.

Like in Section 4, we assume that $\forall k = 1, \dots, K$, the stochastic process X^k is a φ -mixing process, and its mixing coefficients $\varphi^k(i)$ are summable, and we define the following upper bound function \mathcal{M}_k , which differs from the one defined in the previous section:

$$\forall 1 \leq k \leq K, \quad \forall b \in \mathbb{N}, \quad \mathcal{M}_k(b) \doteq 1 + \sum_{i \geq b} \varphi_k(i) < +\infty.$$

In the restless φ -mixing bandit, the agent pulls the arm k in sequences of m_k trials, where the m_k are fixed parameters and may differ for each arm. The mean value of this sequence is defined as follows:

$$\mu^k \doteq \mathbb{E}_{X_1^k, \dots, X_{m_k}^k} \psi_m^k(X_1^k, \dots, X_{m_k}^k)$$

and we use the same definition of regret as defined in Section 3. In the restless setting, an interesting way of dealing with the trade-off exploration/exploitation/independence appears: in addition to the usual exploration, it might be interesting for the agent to pull an apparently sub-optimal arm to get an increased independence on the result of the other arms –since the time between two consecutive sequences of pull decrease their dependency. In order to study this trade-off, we introduce a suitable concentration inequality.

Proposition 3. *Let $1 \leq k \leq K$, $1 \leq m \leq s$, $b = m - s$, $n \in \mathbb{N}^*$, $\psi_m^k : \mathcal{U}^m \rightarrow \mathbb{R}$ be a 1-Lipschitz with respect to the Hamming metric function defined over a countable space \mathcal{U} . Let $\zeta_m^s : \mathbb{N} \mapsto \mathbb{N}$, be such that*

$$\zeta_m^s(t) = t + b \mathbf{1}_{t \geq m+1}.$$

Then the following holds for all $t > 0$:

$$\mathbb{P} \left[\left| \frac{1}{n} \sum_{i=0}^{n-1} \psi_m^k \left(X_{\zeta_m^s(im+1)}^k, \dots, X_{\zeta_m^s((i+1)m)}^k \right) - \mu^k \right| > t \right] \leq 2 \exp \left[-\frac{nt^2}{2\mathcal{M}_k^2(b)} \right]. \quad (17)$$

Proof. This proposition follows from the proof of Theorem 1 using the same technique as Proposition 1. \square

It is interesting to note that the independence trade-off naturally appears in the right term of inequality 17 within the \mathcal{M}_k , and will modify the upper confidence bound. We introduce algorithm 3 to solve this particular setting of φ mixing bandit problem, and we provide a regret analysis for this algorithm.

Proposition 4 (Regret analysis). *Let $R(t)$ be the regret at time t for Algorithm 3. Then, with $\Delta_i = \mu^{k^*} - \mu^i$,*

$$R_t \leq \sum_{1 \leq i \leq K} \left(\Delta_i + \frac{8\alpha \log(t)}{\Delta_i} \right).$$

Algorithm 3 Restless Block-UCB with parameters m_k, α fixed

$t \leftarrow 0, \hat{\psi}_0^k \leftarrow 0, \tau_{k,0} \leftarrow 0, \eta_{k,0} \leftarrow 0, \quad k = 1, \dots, K$
for $t = 1, \dots, \tau$ **do**

Select arm $\hat{k} \in \underset{k}{\operatorname{argmax}} \hat{\psi}_{t-1}^k + \sqrt{\frac{2\alpha \log(t)}{\tau_{k,t-1} \mathcal{M}_k^2(\eta_{k,t})}}$

Update counters and timers :

$$\eta_{\hat{k},t} \leftarrow 0, \tau_{\hat{k},t} \leftarrow \tau_{\hat{k},t-1} + 1,$$

$$\forall k \neq \hat{k}, \quad \tau_{k,t} \leftarrow \tau_{k,t-1}, \text{ and } \eta_{k,t} \leftarrow \eta_{k,t-1} + m_{\hat{k}}$$

Compute the values of the $\hat{\psi}_t^k, m + b \mid s$

$$\hat{\psi}_t^{\hat{k}} \leftarrow \frac{1}{\tau_{\hat{k},t}} \sum_{r=0}^{\tau_{\hat{k},t}-1} \psi_m^{\hat{k}} \left(X_{rm_{\hat{k}}}^{\hat{k}}, \dots, X_{(r+1)m_{\hat{k}}-1}^{\hat{k}} \right)$$

$$\hat{\psi}_t^k \leftarrow \hat{\psi}_{t-1}^k, \text{ for } k \neq \hat{k}$$

end for

Proof. The proof uses the same ideas as the previous regret analysis from Section 3 and 4, and naturally follows from Proposition 3. The proof can be found in the supplementary materials. \square

In this section, we have provided an algorithm for the restless φ -mixing framework. We have seen that the restless setting is a natural framework to use with φ -mixing bandit, as it naturally makes it possible to decrease the dependence of the variables for the arms not chosen. Our algorithm takes advantage of this observation and we were able to show that it has low regret. In order to do so, in addition to the usual τ_k , the number of time a given arm has been pulled, the Algorithm 3 computes the η_k , the time spent since the last time the arm k was sampled. Indeed, as seen in (17), as η_k increases, $\mathcal{M}_k^2(\eta_k)$ decreases and the optimistic value of the arm increases.

6 Conclusion

We have studied an extension of the multi-armed bandit problem to the stationary φ -mixing framework, both in the rested and in the restless case. We have provided both a theoretical analysis in a general framework, and a more practical study of the problem in the case of fast mixing sequences (with $\sum \varphi(i) < +\infty$). For each of these cases, we provided algorithms and accompanying regret analyses, which are strict extensions of the methods that exist for the i.i.d situation, as usual results might be recovered from our bounds when the mixing coefficients are all 0. Future works might include a study of the restless case where the m_k has to be computed from the data, as well as a study in the more difficult case of β -mixing processes.

References

- [1] T. L. Lai and H. Robbins. Asymptotically efficient adaptive allocation rules. *Advances in Applied Mathematics*, 6:422, 1985.
- [2] P. Auer, N. Cesa-Bianchi, and P. Fischer. Finite-time analysis of the multi-armed bandit problem. *Machine Learning Journal*, 47(23):235–256, 2002.
- [3] P. Auer and R. Ortner. Ucb revisited: Improved regret bounds for the stochastic multi-armed bandit problem. *Periodica Mathematica Hungarica*, 61:5565, 2010.
- [4] S. Bubeck and N. Cesa-Bianchi. *Regret Analysis of Stochastic and Nonstochastic Multi-armed Bandit Problems*, volume 5 of *Foundation and Trends in Machine Learning*. NOW, 2012.
- [5] R. Ortner, D. Ryabko, P. Auer, and R. Munos. Regret bounds for restless markov bandits. Elsevier, (to appear).
- [6] C. Tekin and M. Liu. Online learning of rested and restless bandits. *IEEE Transactions on Information Theory*, 58(8):5588–5611, 2012.
- [7] L. Kontorovich and K. Ramanan. Concentration inequalities for dependent random variables via the martingale method. *The Annals of Probability*, 36(6):2126–2158, 2008.
- [8] M. Mohri and A. Rostamizadeh. Stability bounds for stationary α -mixing and β -mixing processes. *Journal of Machine Learning Research*, 11:789–814, 2010.

A Appendix

Proposition 5. Let $m, b \in \mathbb{N}$ and $s = m + b$, X_t be a φ -mixing process on Ω taking value in \mathbb{R} , with mixing coefficient $\varphi_X(\cdot)$, and $\psi : \mathbb{R}^m \mapsto \mathbb{R}$ be a measurable function. Then the stochastic process Z_t defined by

$$Z_t = \psi(X_{st+1}, \dots, X_{st+m})$$

is also a φ -mixing process with mixing coefficient $\varphi_Z = \varphi_X \circ \kappa$, where $\kappa(t) = bt + m(t-1)$.

Proof. In the following we use $\sigma(A)$ to denote the σ -algebra generated by A . First note that since ψ is measurable $\sigma(\psi^{-1}(\mathbb{R})) \subset \sigma(\mathbb{R}^m)$, and as a consequence

$$\sigma(Z_t) \subset \sigma(X_{st+1}, \dots, X_{st+m})$$

(since σ -algebra are closed under countable intersection).

Now for any $i, j \in \mathbb{Z} \cup \{-\infty, +\infty\}$, let $\sigma_i^j(Z)$ denote the σ -algebra generated by the random variables Z_k , $i \leq k \leq j$. Then, for any positive integer k , the φ -mixing coefficient $\varphi_Z(t)$ of the stochastic process \mathbf{Z} is defined as

$$\begin{aligned} \varphi_Z(t) &= \sup_{n, A \in \sigma_{n+t}^{+\infty}(Z), B \in \sigma_{-\infty}^n(Z)} |\mathbb{P}[A|B] - \mathbb{P}[A]| \\ &\leq \sup_{n, A \in \sigma_{s(n+t)}^{+\infty}(X), B \in \sigma_{-\infty}^{ns+m}(X)} |\mathbb{P}[A|B] - \mathbb{P}[A]| \\ &= \varphi_X(s(n+t) - ns - m) = \varphi_X(tb + (t-1)m) \end{aligned}$$

□

B Proof of Proposition 2

PROOF :

If at time t , the arm chosen is k instead of k^* , then one of the following must be true :

$$\psi_t^{k^*} \leq \nu_{m^*, b^*}^{k^*} \tag{18}$$

$$\exists m + b \mid s, \quad \nu_{m^*, b^*}^k \leq \beta_{m, b} \left(\hat{\psi}_{m, b, t-1}^k - \sqrt{\alpha} J_{m, b}(k, t-1) \right) \tag{19}$$

$$\exists m + b \mid s, \quad 2\sqrt{\frac{2\alpha\beta_{m, b} \log(t)}{\tau_{k, t-1}}} \geq \nu_{m^*, b^*}^{k^*} - \nu_{m^*, b^*}^k \tag{20}$$

Otherwise, $\forall m + b \mid s$,

$$\begin{aligned} \psi_t^{k^*} &\geq \nu_{m^*, b^*}^{k^*} \geq \nu_{m^*, b^*}^k + 2\sqrt{\frac{2\alpha\beta_{m, b} \log(t)}{\tau_{k, t-1}}} \\ &\geq \nu_{m^*, b^*}^k + 2\sqrt{\frac{2\alpha\beta_{m, b} \log(t)}{\tau_{k, t-1}}} \\ &\geq \beta_{m, b} \left(\hat{\psi}_{m, b, t-1}^k + \sqrt{\alpha} J_{m, b}(k, t-1) \right) \end{aligned}$$

Since the last line is true $\forall m + b \mid s$, we deduce that $\psi_t^{k^*} \geq \psi_t^k$ hence $T(t) \neq k$, which is absurd. Then, we need to bound the probability of the events defined by (18), (19) and (20).

Because $\psi_t^{k^*}$ is defined as a maximum, we have

$$\begin{aligned} \mathbb{P}(\psi_t^{k^*} \leq \nu_{m^*, b^*}^{k^*}) &\leq \mathbb{P}\left(\beta_{m^*, b^*}(\hat{\psi}_{m^*, b^*, t-1}^{k^*} - \sqrt{\alpha} J_{m^*, b^*}(k^*, t-1)) \leq \nu_{m^*, b^*}^{k^*}\right) \\ &\leq \frac{1}{t^\alpha}, \end{aligned}$$

using Proposition 1.

Now, by definition of $\nu_{m_k^*, b_k^*}^k$,

$$\begin{aligned}
& \mathbb{P} \left(\exists m + b \mid s, \quad \nu_{m_k^*, b_k^*}^k \leq \beta_{m,b}(\widehat{\psi}_{m,b,t-1}^k - \sqrt{\alpha} J_{m,b}(k, t-1)) \right) \\
& \leq \mathbb{P} \left(\exists m + b \mid s, \quad \nu_{m,b}^k \leq \beta_{m,b}(\widehat{\psi}_{m,b,t-1}^k - \sqrt{\alpha} J_{m,b}(k, t-1)) \right) \\
& \leq \sigma(s) \max_{m+b \mid s} \mathbb{P} \left(\nu_{m,b}^k \leq \beta_{m,b}(\widehat{\psi}_{m,b,t-1}^k - \sqrt{\alpha} J_{m,b}(k, t-1)) \right) \\
& \leq \frac{\sigma(s)}{t^\alpha}
\end{aligned}$$

where we used Proposition 1 again at the last line.

Finally,

$$\begin{aligned}
& \left\{ \exists m + b \mid s, \quad 2\sqrt{\frac{2\alpha\beta_{m,b}\log(t)}{\tau_{k,t-1}}} \geq \nu_{m^*, b^*}^{k^*} - \nu_{m_k^*, b_k^*}^k \right\} \subset \left\{ 2\sqrt{\frac{2\alpha s \log(t)}{\tau_{k,t-1}}} \geq \nu_{m^*, b^*}^{k^*} - \nu_{m_k^*, b_k^*}^k \right\} \\
& = \left\{ \frac{8\alpha s \log(t)}{(\nu_{m^*, b^*}^{k^*} - \nu_{m_k^*, b_k^*}^k)^2} \geq \tau_{k,t-1} \right\}
\end{aligned}$$

i.e. the event defined by (20) happens at most $u = \left\lceil \frac{8\alpha s \log(t)}{(\nu_{m^*, b^*}^{k^*} - \nu_{m_k^*, b_k^*}^k)^2} \right\rceil$ times.

hence the conclusion □

C Proof of Proposition 5

Proof. If at time t , the arm chosen is k instead of k^* , then one of the following must be true :

$$\psi_t^{k^*} \leq \mu^{k^*} \tag{21}$$

$$\mu^k \leq \left(\widehat{\psi}_{t-1}^k - \sqrt{\alpha} J(k, t-1) \right) \tag{22}$$

$$2\sqrt{\frac{2\alpha \log(t)}{\mathcal{M}_k^2(\eta_{k,t})\tau_{k,t-1}}} \geq \mu^{k^*} - \mu^k \tag{23}$$

From the last inequality, we deduce that :

$$\tau_{k,t-1} \leq \frac{8\alpha \log(t)}{\mathcal{M}_k^2(\eta_{k,t})\Delta_k^2} \leq \frac{8\alpha \log(t)}{\Delta_k^2} \tag{24}$$

since $\mathcal{M}_k^2(\eta_{k,t}) \leq 1$. □